(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
7 November 2002 (07.11.2002)

PCT

(10) International Publication Number
WO 02/089057 A1

(51) International Patent Classification⁷: G06T 1/00, H04N 1/32, 7/26, 7/24

(21) International Application Number: PCT/GB02/01849

(22) International Filing Date: 23 April 2002 (23.04.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0110150.0 25 April 2001 (25.04.2001) GB

(71) Applicant *(for all designated States except US)*: **UNIVERSITY OF BRISTOL** [GB/GB]; 3rd floor, Senate House, Tyndall Avenue, Bristol BS8 1TH (GB).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **BULL, David, Roger** [GB/GB]; Netherways, Netherhope Lane, Tidenham, Near Chepstow, Monmouthshire NP6 7JE (GB). **CANAGARAH, Cedric, Nishan** [GB/GB]; 29 Elberton Road, Coomble Dingle, Bristol BS9 2PZ (GB). **WINNE, Dominic** [BE/GB]; 25 Fanshawe Road, Hengrove, Bristol BS14 9RY (GB). **KNOWLES, Henry, David** [GB/GB]; 23 Fernbank Road, Bristol BS6 6PZ (GB).

(74) Agent: **O'CONNELL, David, Christopher**; Haseltine Lake & Co., Imperial House, 15-19 Kingsway, London WC2B 6UD (GB).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,

[Continued on next page]

(54) Title: WATERMARKING USING REPRESENTATIVE VALUES

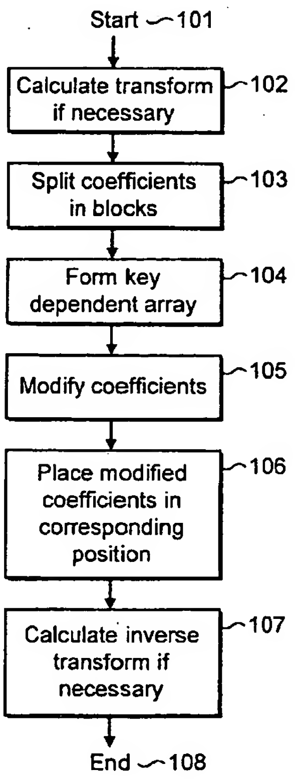
Start 101
Calculate transform if necessary
102
Split coefficients in blocks
103
Form key dependent array
104
Modify coefficients
105
Place modified coefficients in corresponding position
106
Calculate inverse transform if necessary
107
End 108

(57) Abstract: The present invention provides a scheme for inserting and extracting an invisible watermark in a image/video sequence, in order to verify whether a watermarked image has been tampered with. According to one aspect of the invention, there is provided a method of watermarking a digital image, comprising: obtaining digital image coefficients; forming a plurality of arrays of image coefficients, each array comprising a predetermined number of image coefficients; selecting a plurality of input values, selected input values each representing binary 0s or 1s; and for each array: determining a binary value as a watermark component bit; calculating a numerical value being representative of the image coefficients; and, if necessary, modifying one or more of the image coefficients, such that the representative numerical value becomes equal to a selected input value which represents the determined watermark component bit. According to other aspects of the invention, there are provided a corresponding method of detecting a watermark, and corresponding hardware and software devices.

WO 02/089057 A1

SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) **Designated States** *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

WATERMARKING USING REPRESENTATIVE VALUES

This invention relates to watermarking, and in particular to a system for embedding and extracting a fragile watermark in a digital object, such as a digital image.

More specifically, the invention relates to techniques for inserting and extracting fragile watermarks in an image, for the purposes of image authentication, allowing a user to determine whether and how a watermarked image has been tampered with.

A watermark is a visible or invisible structure in an image, which can be recovered after it has been embedded. A digital watermark is a digital pattern inserted into a digital creation, such as a digital image. The process of inserting a watermark into a digital image (embedding procedure) can be done directly in the spatial or transformed domain. The watermark can be inserted by altering certain coefficients in a way which minimises the resulting distortion of the image. The imperceptibility of the watermark is the first line of defence, since, if an image is not visibly watermarked, it is more difficult to avoid the watermark, by tampering with the image undetectably.

Most watermarking techniques fall into one of two main categories, namely, the robust type, for copyright protection, and the fragile type, for authentication applications.

Robust watermarks must be embedded in the image so that it is virtually impossible, or at least difficult, to remove the watermark without visibly damaging the image. Such a watermark must be resistant to several image-processing techniques, such as cropping, scaling, filtering, compression/decompression, etc.

Fragile watermarks are designed to detect any unauthorised alteration of the image signal. Fragile watermarks may be used, for example, in connection with images generated by digital surveillance cameras to provide a basis for determining whether or not an image has been tampered with after its creation. The requirements of a fragile watermarking system change, depending on the data type and application.

An effective authentication scheme should have some or all of the following desirable features:

to be able to determine whether an image has been altered or not, and to be able to locate any alteration made on the image, even without having the original image data;

to be able to integrate the authentication data with host image data, rather than as a separate data file;

the embedded authentication data should be invisible under normal viewing conditions;

to be able to extract information concerning the method of attack;

to be able to restore the attacked area in the watermarked image.

US-6,064,764 proposes a fragile digital watermark embedding and extraction system where the insertion process is accomplished by embedding the bits of a digital signature of a hash function of the image in the frequency coefficients of the image. This procedure results in an imperceptible watermark, but the properties of hash functions prevent the localisation or quality assessment of the changes that have been made to the image.

EP-A-0901102 describes embedding watermarking information in respective pixels or by using a relation between pixels or as a position pattern of changed pixels in the contents. This robust embedding system in the spatial domain uses a similar embedding function where a binary watermark value changes a specific number of a block of pixels.

EP-A-0953938 describes an invisible digital watermark insertion technique for ownership verification or/and authentication purposes. The watermark is a combination of a given watermark bitmap and the hashed image. The extreme localisation is accomplished by embedding the digital signature in the Least Significant Bit (LSB) band of the image. However the method cannot distinguish malicious changes from innocent image processing operations.

The present invention seeks to provide a scheme for inserting and extracting an invisible fragile watermark in an image/video sequence, in order to verify whether a watermarked image has been tampered with.

According to one aspect of the invention, there is provided a method of watermarking a digital image, comprising:

obtaining digital image coefficients;

forming a plurality of arrays of image coefficients, each array comprising a predetermined number of image coefficients;

selecting a plurality of input values, selected input values each representing binary 0s or 1s; and

for each array:

determining a binary value as a watermark component bit;

calculating a numerical value from the image coefficients in the array, the calculated numerical value being representative of the image coefficients; and,

if necessary, modifying one or more of the image coefficients, such that the representative numerical value becomes equal to a selected input value which represents the determined watermark component bit.

According to other aspects of the invention, there are provided a corresponding method of detecting a watermark, and corresponding hardware and software devices.

For a better understanding of the invention, and to show how it may be put into effect, reference will now be made to the accompanying drawings.

Figure 1 is a flow chart illustrating a method of embedding a watermark, in accordance with the invention.

Figure 2 illustrates a step in the method of Figure 1.

Figure 3-5 illustrate further steps in the method of Figure 1.

Figure 6 further illustrates the method of Figure 1.

Figure 7 is a flow chart illustrating in more detail the method of embedding a watermark.

Figure 8 is a flow chart illustrating a method of extracting a watermark, in accordance with the invention.

Figure 9 illustrates a part of the extraction process.

Figures 10 and 11 illustrate a method of optimising the extraction process.

Figure 12 illustrates the result of the extraction optimisation.

Figure 13 illustrates a further method of embedding a watermark, in accordance with the invention.

Firstly, Figure 1 illustrates an embodiment of the invention in which a watermark is embedded in an image.

In step 101, the procedure starts with a set of image data. The image data may be in any form. For example, the image data may be a set of luminance and chrominance values associated with respective pixels of an image. The embedding of the watermark is advantageously carried out close to the source of the image. For example, the watermarking procedure can be carried out within a digital camera, such as a surveillance camera.

The watermark embedding process can be carried out by a general purpose computer, operating under the control of suitable software, or by another hardware device, such as a DSP or an ASIC, or other integrated circuit.

The method of embedding the watermark into a digital image signal can be carried out in the spatial or transformed domain, depending on the application. Thus, step 102 shows a transformation procedure. . The transformation may be of any type, and may for example involve performing a digital cosine transform (DCT), a Fourier transform, a DFT, DST, Walsh, Hadamard, Hartley, or wavelet transform. Thus, when referring herein to coefficients, these can represent pixel values for a greyscale or coloured image, or DCT, wavelet or any transformed coefficients of the digital signal.

Which coefficients will be selectable by the user, for example depending on the application of the technique. For example, in the case of a DCT transformed block, it may be preferable not to modify the DC coefficient, or in the case of a wavelet transformed image it may be preferable to modify only the low-pass coefficients, or it may be preferable to allow only the largest value coefficients to be modified.

Next, in step 103, the coefficients of the image are divided into blocks. After watermarking, each block will contain a watermark bit value, and so the block size represents the detection resolution. The introduced modification is spread over the selected coefficients. Therefore, the smaller the block size, the better the detection resolution, but the greater the distortion of the coefficients, and hence the larger the visual degradation of the watermarked image.

In step 104, a key dependent array is formed, in order to hide the watermark "skeleton", as shown in more detail in Figure 2.

Thus, as shown in Figure 2, the blocks of coefficients 201 are divided into sub-blocks 202, and an array 203 constructed from each sub-block in a zig-zag order. A random generator 204, controlled by a user key 205, then uniquely creates an array 206 of coefficients which can be further processed.

Without access to the secret user key 205, a person attempting to tamper with the image would not know which pixels corresponded to any given block. The secret key acts as a seed number for a random number generator. The random number which is generated then controls the way in which the arrays of coefficients are formed.

The random number also provides a secret key which is used in the same way in the next block, and so on. This means that the watermark is effectively embedded in different blocks in different ways. This provides a defence against an attempt to tamper with an image by forming a collage of blocks from different images or from different parts of the same image.

The initial key, which governs the creation of the secret keys used in the respective blocks, is stored in the device in which the watermark is embedded, such that it can be retrieved only by an authorised person.

The formation of the arrays in this way means that, if tampering of an image is detected, it is possible to determine the region of the image which has been tampered with.

An alternative way of achieving this is to form at least two, and preferably more than two, mutually orthogonal sets of arrays, with the arrays each being formed from coefficients selected from anywhere in the image, and with one watermark bit being embedded in each array. At the detection stage, any tampering of the image will result in a detection in each of the sets of arrays, thereby allowing localisation of the tampering.

Once the arrays have been formed, then, in step 105 of Figure 1, the watermark bit is embedded in the selected array by modifying the values of the coefficients in each array, according to an algorithm to be described below. In steps 106 and 107, the watermarked image is obtained, by replacing the modified coefficients in their original position and then by calculating the inverse of any transform performed in step 102.

The basis of the algorithm which is used is that, for each array of coefficients, a representative value is calculated. The coefficient values are then manipulated in such a way that the representative value takes a modified value. More specifically, one group of representative values are preselected to represent binary 0s, and another group of representative values are preselected to represent binary 1s, the values which represent binary 0s and 1s being referred to as input values. Then, when a watermark bit is applied to each array, the coefficient values are manipulated such that the representative value takes a modified value which is equal to an input value corresponding to the watermark bit.

To illustrate this, reference will now be made to Figure 3, which shows how to calculate a representative value for the illustrated series of digits 10, 5, 12, 6, 6, 9, 2, 4, 11, 1.

Specifically, the representative value is the sum of the digits in odd-numbered positions (i.e. 41) minus the sum of the digits in even-numbered positions (i.e. 25). Thus, the representative value is 16 in this case.

Where the sum of the digits in odd-numbered positions is greater than the sum of the digits in even-numbered positions, the array is referred to as a positive array. In the opposite case, the array is referred to as a negative array, and the absolute value of the calculated value is taken as the representative value.

As mentioned above, the algorithm of the preferred embodiment of the invention involves manipulating the coefficients of the array, so that the representative value of the array takes a value which is one of a small number of available values.

Figures 4 and 5 show an example of a manipulation procedure, purely for illustrative purposes.

Thus, in Figure 4, the series of digits shown in Figure 3, which formed a positive array with a representative value of 16, is manipulated such that the representative value becomes equal to 10. Specifically, the digits in the first three odd-numbered positions are each decreased by one, while the digits in the first three even-numbered positions are each increased by one.

Figure 5 shows an alternative series of digits, forming a negative array which also had a representative value of 16. The series is again manipulated such that the representative value becomes equal to 10. In this case, the digits in the first three odd-numbered positions are each increased by one, while the digits in the first three even-numbered positions are each decreased by one.

In the preferred embodiment of the present invention, input values [0, 200, 400, 600, ...] are chosen to indicate binary 0s, and input values [100, 300, 500, 700, ...] are chosen to indicate binary 1s, as shown in Figure 6. Then, for each array, when the representative value has been calculated, and the watermark bit has been determined, the coefficient values within the array are modified such that the modified representative value becomes equal to the closest input value which indicates the binary value of the watermark bit.

Increased separation of adjacent input values (100 in this illustrated example) improves the detector performance, but introduces more distortion in the watermarked digital signal.

Figure 7 illustrates the procedure for modifying the coefficient values in a positive array, in order to manipulate the representative value to become equal to one of the input values. In the case of a negative array, blocks 704 and 708 must be reversed.

Thus, the procedure starts in step 701 by generating the required closest input value which, as described above, depends on the representative value, the watermark bit to be embedded, and the user defined input values.

Then, step 702, a counter i, which indicates the specific one of the coefficients C_i in the array that is under consideration, is initialised. Next, in step 703, it is detected whether the specific representative value calculated for the array is equal to one of the input values. If so, then the modification is finished for this array, but, if not, the coefficient C_i that is under consideration will be changed.

For example, consider the case of an array where the watermark bit which is to be embedded, is a 1, and the array is a positive array, with a calculated specific representative value of 295.

From Figure 6, it can be seen that the input value, which is closest to the calculated representative value, and which represents a binary 1 is 300.

In order to manipulate the coefficient values, such that the representative value become equal to 300, while minimising the introduced distortion, five of the coefficients in the array are modified by 1.

Thus, this describes a system where each coefficient is modified by the same amount, as nearly as possible. However, particularly in cases where a greater degree of modification of the coefficients is required, it may be advantageous to modify the coefficients by amounts which are approximately proportional to the values of the coefficients, or to make modifications which introduce the smallest amount of distortion into the image.

As shown in Figure 7, 1 is added to the coefficient on the i-th location in step 704. In step 705, the counter i is increased until the end of the array is reached, as tested in step 706. After each modification, the specific value is recalculated, to detect whether the array of coefficients has been sufficiently modified to obtain the desired input value.

Then, in step 708, 1 is subtracted from the following coefficient and again the counter i is incremented in step 709.

Thus, this describes a system whereby the modifications are applied preferentially to the first coefficients in the array. However, it may be preferable to modify the coefficients in a different order, for example by changing the way in which the counter is incremented.

In the case of a negative array, a modified procedure is used, in which 1 is subtracted from the selected coefficient value in step 704 and added to the selected coefficient value in step 708.

Figure 8 is a flow chart which illustrates the process for extraction and detection of a watermark from a watermarked image. This relies on the knowledge that, if the image coefficients are divided into arrays, in the same way that an image is divided to allow insertion of a watermark, then, if the watermark remains intact, each of the calculated specific representative values of the array will be equal to one of the input values. Moreover, the pattern of binary values represented by these input values will represent the watermark pattern.

If the image has been modified to any significant extent, it is extremely unlikely that the extracted watermark will equal to the embedded one.

The extraction process of Figure 8 first calculates the transform coefficients in step 802, if any transform is used, and then groups the coefficients together in arrays in step 803, and generates the key dependent array in step 804, in the same way as during the embedding process, as shown in Figure 2.

Then, in step 805, the specific number of the selected array is calculated, and it is determined whether this lies within a selected tolerance band. The use of the tolerance bands will now be described with reference to Figure 9.

Figure 9 shows an example based on the selected input values shown in Figure 6, in which embedded watermark bits having the value 0 are used to modify the specific value of an array to 0 etc, while embedded watermark bits having the value 1 are used to modify the specific value of an array to 100 etc.

In the case of Figure 9, calculated specific values in the tolerance band 0±10 (say) are used to infer that a 0 has been embedded, while calculated specific values in the tolerance band 100±10 (say) are used to infer that a 1 has been embedded. This produces a non-dichotomous detector, in which there is a band between the two tolerance bands, which always leads to a determination that the image coefficients must have been modified after the watermark has been embedded.

The use of tolerance bands is necessary because there is always the possibility of some distortion as a result of legitimate image processing.

Alternatively, the tolerance bands can be set so that values in the tolerance band 0±50 are used to infer that a 0 has been embedded, while calculated specific values in the tolerance band 100±50 are used to infer that a 1 has been embedded. This produces a dichotomous detector.

Thus, based on the calculated specific value, it may be possible to extract a binary value which represents this array of coefficients.

If the calculated specific value does not lie within one of the tolerance bands, the process passes from step 805 to step 807, and it is determined that the extent of modification of the coefficients means that this area of the image must have been modified.

If the calculated specific value does lie within one of the tolerance bands, and the resulting extracted binary value is equal to the embedded watermark bit, the process passes from step 805 to step 806, to step 809 and to step 810, and it is determined that the watermark bit has been successfully detected, meaning that there is no evidence of tampering with this area of the image.

If the calculated specific value does lie within one of the tolerance bands, but the resulting extracted binary value is opposite to the embedded watermark bit, the process passes from step 805 to step 806, to step

809 and to step 807, and it is determined that the extent of modification of the coefficients means that this area of the image must have been modified.

It is advantageous to be able to refine the detector output, in order to be able to detect tampering of images more reliably. That is, the detector output for an array, obtained from the procedure described above, may be changed, for consistency with the detector outputs obtained for other nearby arrays.

Figures 10(a) and 10(b) show two situations where white squares 150 represent blocks where the detector output indicates no tampering, shaded squares 160 represent detector outputs which indicate tampering, and the square 170 is under consideration. In both cases, the detector output gives no evidence of tampering. However, in Figure 10(a) the square 170 is within an area which has not been tampered with, while in Figure 10(b) the square 170 is within an area which has been tampered with.

In this illustrated embodiment of a detector optimisation function, a weighted function is used to distinguish between these situations, and hence determine the refined or optimised detector output for an array, based on the outputs for neighbouring blocks.

In Figure 10(a) and Figure 10(b), three of the blocks neighbouring the square 170 have been attacked in each case.

The weighting function distinguishes between the two cases, as shown in Figure 11. Specifically, when considering a particular block 180, marked with an "X", the detector outputs for the eight neighbouring blocks 190 are considered, and given the weightings "1" or "3", shown in Figure 11. It will be appreciated that this arrangement of weightings is purely exemplary.

Then, if neighbouring blocks with weightings which are equal to or higher than a threshold value of 6, the detection output for the block under consideration may be changed from an output which indicates no evidence of tampering, to an output which indicates tampering.

The nature of the process is such that the detection process is more likely to produce a false negative result than a false positive. Therefore, the optimization process can be carried out only on blocks for which the unrefined detection output indicates no evidence of tampering.

Alternatively, blocks for which the unrefined detection output indicates no evidence of tampering can be tested against one threshold value, while blocks for which the unrefined detection output does indicate tampering can be tested against a second, higher threshold value.

Thus, returning to Figure 10(a), three of the eight blocks neighbouring the square 170 indicate an attack, and these three blocks have weightings 1, 3 and 1, totalling 5. Since this is less than 6, it is determined that the unrefined detector output for the block, indicating no tampering, should not be changed.

In Figure 10(b), again, three of the eight blocks neighbouring the square 170 indicate an attack, but in this case these three blocks have weightings 3, 1 and 3, totalling 7. Since this is greater than 6, the refinement function determines that the initial detector output for the block, indicating no tampering, should be changed.

Figure 12 shows the result of applying this detector optimisation. Figure 12(a) illustrates a non-optimised detector output, in which black squares indicate blocks which have been determined to have been tampered with, while white squares indicate blocks which have in fact been tampered with, although the initial detector outputs have not detected evidence of tampering.

Since there is a consistent pattern of tampering, it is probable that these provisional determinations, that there has not been tampering, are false negative results. Indeed, the refinement process, described above, is able to determine that these results should be changed.

Thus, Figure 12(b) illustrates the optimised detector output, in which most of these provisional determinations of non-tampering have been modified.

Thus, if any part of the image is changed, the watermark extraction procedure, according to the present invention, will return an output that indicates that specific parts of the image that have been changed. If a watermarked image is resized or cropped, then the extraction procedure will also return an output that indicates that the whole image was changed.

One implementation of this invention is in a system, such as the JPEG image encoding system, which quantizes coefficients. The effect of quantization is to change coefficient values to the nearest allowed quantized level. In accordance with the invention, the watermark can be embedded by modifying a coefficient value before it is quantized, but with knowledge of what the result of the quantization would have been.

Figure 13 illustrates a further method of embedding a watermark, in accordance with the invention.

In step 901, the image coefficients are divided into blocks, for example 8x8 blocks. In step 902, the coefficients within each block are permuted in accordance with an algorithm. In step 903, any required transform is applied. In step 904, the watermark is embedded, as described above, for example in steps 104-106 of Figure 1. In step 905, the inverse of any transform applied in step 903 is applied. In step 906 the inverse of the permutation applied in step 902 is applied, and then, in step 907, the watermarked image is obtained.

There is therefore described a watermarking system which can be used in many different systems. In the case of colour images, the watermark can be embedded in every plane or just a specific plane.

In the case of a video sequence, every frame, or just some selected frames, e.g. I-frames, can be watermarked. The choice of which frames should be watermarked depends on the application and the ability to detect frame reordering and attack by frame dropping or frame removal.

The strength of the verification lies in the uniqueness of the key which is used.. Any cryptographic algorithm can generate a key that would enable the correct-key-holding user to detect the watermark by mixing the coefficients in the unique order before executing the detection process. The secret key could also control the division of the image in blocks, or the permutation of the coefficients, or could contain information about the watermark.

CLAIMS

1. A method of watermarking a digital image, comprising:

obtaining digital image coefficients;

forming a plurality of arrays of image coefficients, each array comprising a predetermined number of image coefficients;

selecting a plurality of input values, selected input values each representing binary 0s or 1s; and

for each array:

determining a binary value as a watermark component bit;

calculating a numerical value from the image coefficients in the array, the calculated numerical value being representative of the image coefficients; and,

if necessary, modifying one or more of the image coefficients, such that the representative numerical value becomes equal to a selected input value which represents the determined watermark component bit.

2. A method as claimed in claim 1, wherein the arrays of image coefficients are selected in accordance with a cryptographic key.

3. A method as claimed in claim 1 or 2, wherein the image coefficients comprise transform coefficients.

4. A method as claimed in claim 3, wherein the image coefficients comprise quantized coefficients.

5. A method as claimed in claim 1 or 2, wherein the image coefficients comprise bit values.

6. A computer system, programmed to carry out a method as claimed in any one of claims 1 to 5.

7. An image generation device, comprising:

means for generating a digital image; and

a computer system containing software for carrying out a method as claimed in any one of claims 1 to 5.

8. A computer software product, containing code for carrying out a method as claimed in any one of claims 1 to 5.

9. A hardware device, adapted to carry out a method as claimed in any one of claims 1 to 5.

10. A method of detecting a watermark in a digital image, the method comprising:

obtaining digital image coefficients;

forming a plurality of arrays of image coefficients, each array comprising a predetermined number of image coefficients, the image coefficients;

selecting a plurality of input values, selected input values representing binary 0s or 1s, each selected input value having a tolerance band associated therewith; and

for each array:

determining a binary value as an actual watermark component bit;

calculating a numerical value from the image coefficients in the array, the calculated numerical value being representative of the image coefficients; and

determining a detected watermark component bit on the basis of a determination that the representative numerical value lies within a tolerance band associated with a resulting input value, and on the basis of the binary value associated with said resulting input value; and

comparing the detected watermark component bit and the actual watermark component bit.

11. A method as claimed in claim 10, further comprising:

displaying said digital image, distinguishing between regions of the image in which the detected watermark component bit matches the actual watermark component bit, and regions of the image in which the detected watermark component bit does not match the actual watermark component bit.

12. A method as claimed in claim 10 or 11, further comprising, if the detected watermark component bit matches the actual watermark component bit, determining whether to modify a detection output, on the basis of a comparison result obtained in one or more adjacent arrays.

13. A computer system, programmed to carry out a method as claimed in any one of claims 10-12.

14. A computer software product, containing code for carrying out a method as claimed in any one of claims 10-12.

15. A hardware device, adapted to carry out a method as claimed in any one of claims 10-12.

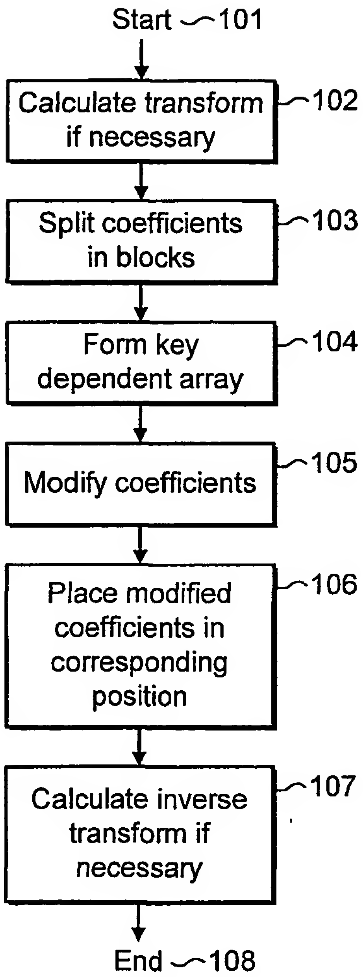
Start 101
Calculate transform if necessary 102
Split coefficients in blocks 103
Form key dependent array 104
Modify coefficients 105
Place modified coefficients in corresponding position 106
Calculate inverse transform if necessary 107
End 108

FIG. 1

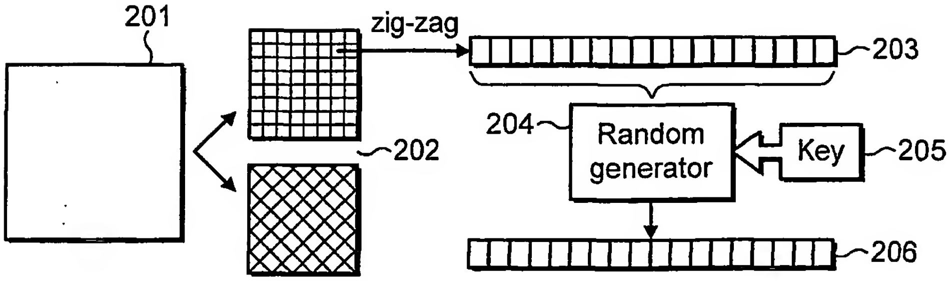
201
zig-zag
203
202
204
Random generator
Key
205
206

FIG. 2

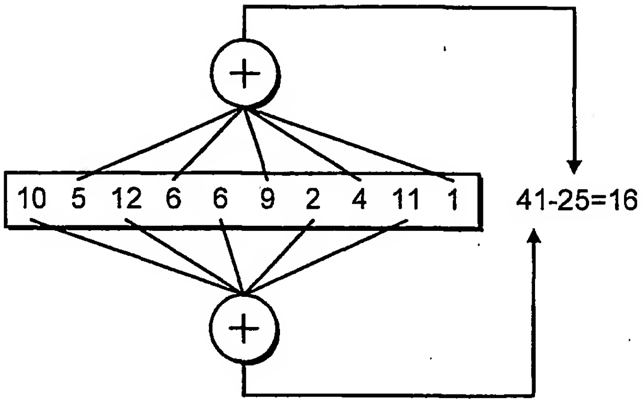

FIG. 3

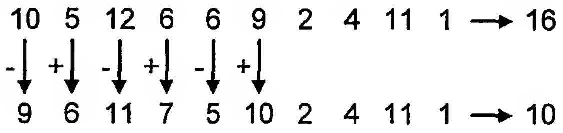

FIG. 4

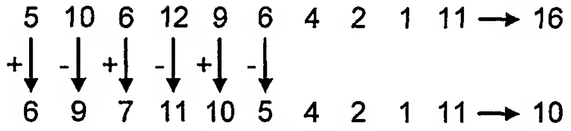

FIG. 5

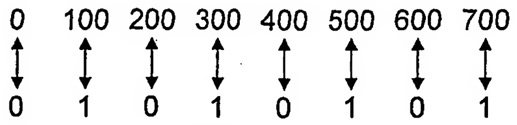

FIG. 6

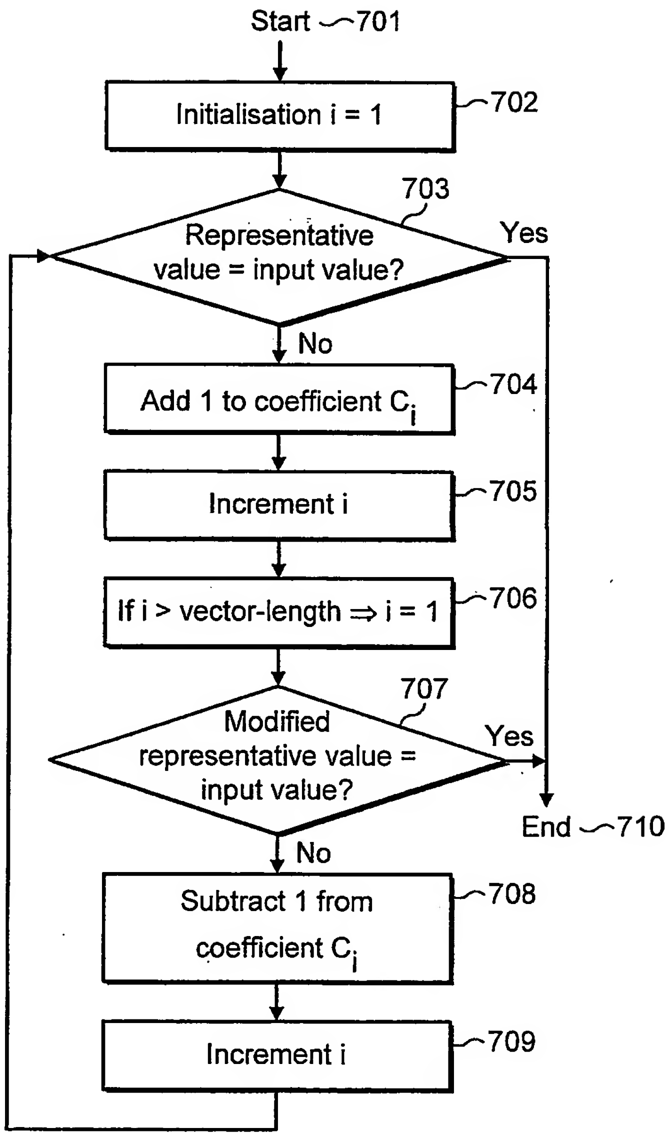
Start 701
Initialisation i = 1
702
703
Representative value = input value?
Yes
No
Add 1 to coefficient C_i
704
Increment i
705
If i > vector-length ⇒ i = 1
706
707
Modified representative value = input value?
Yes
End 710
No
Subtract 1 from coefficient C_i
708
Increment i
709

FIG. 7

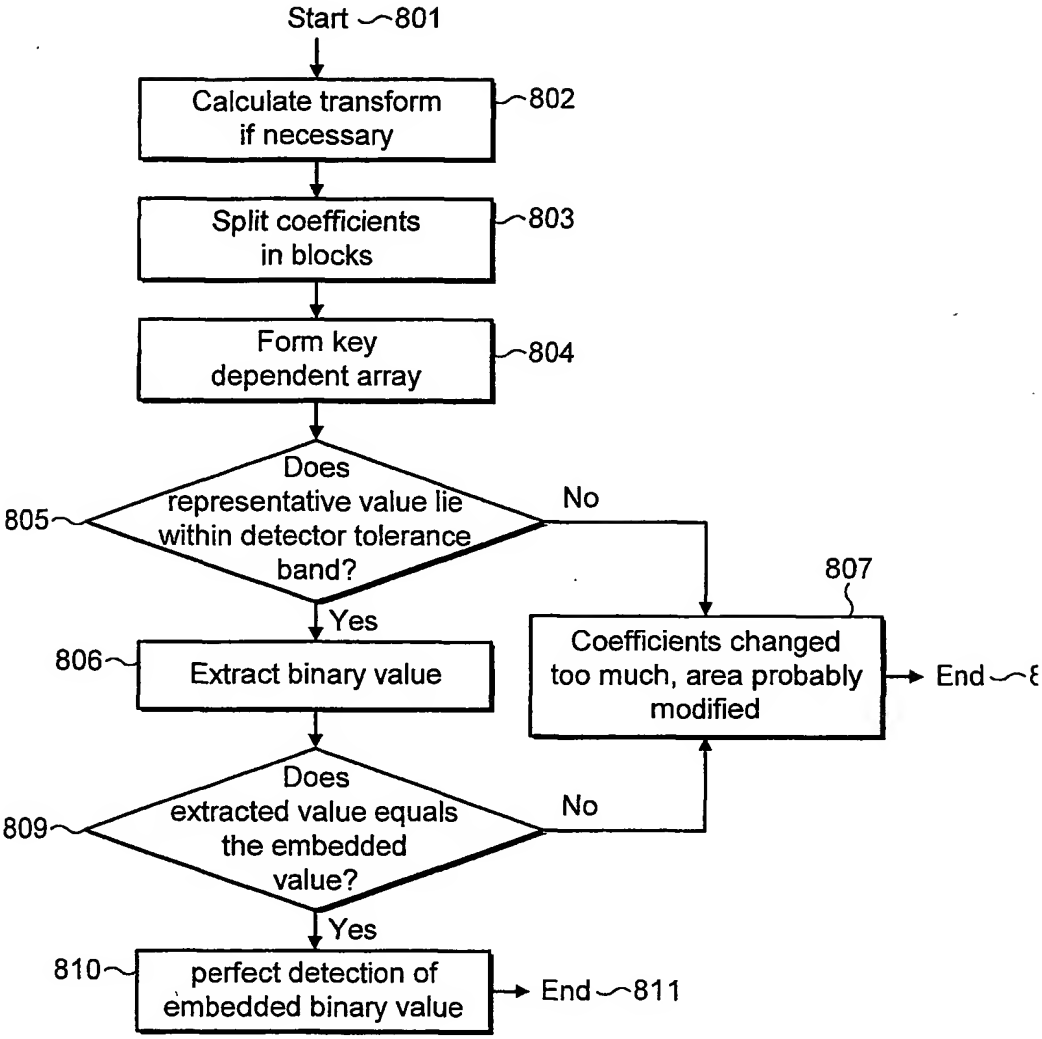
Start 801
Calculate transform if necessary
802
Split coefficients in blocks
803
Form key dependent array
804
805
Does representative value lie within detector tolerance band?
No
Yes
807
806
Extract binary value
Coefficients changed too much, area probably modified
End 8
809
Does extracted value equals the embedded value?
No
Yes
810
perfect detection of embedded binary value
End 811

FIG. 8

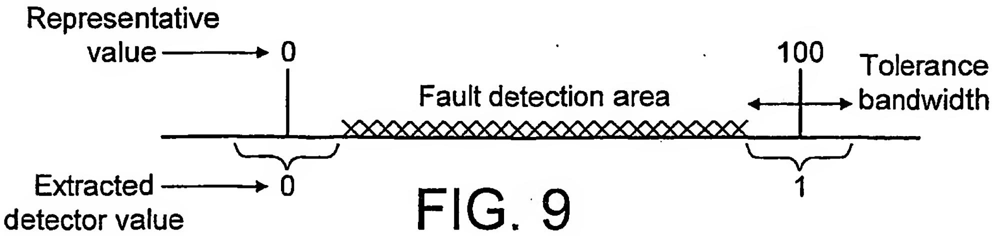
Representative value
0
100
Tolerance bandwidth
Fault detection area
Extracted detector value
0
1

FIG. 9

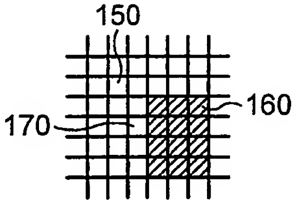

FIG. 10(a)

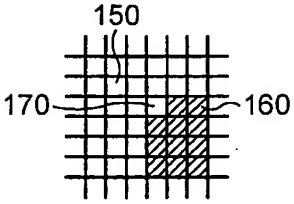

FIG. 10(b)

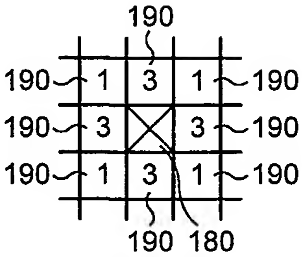

FIG. 11

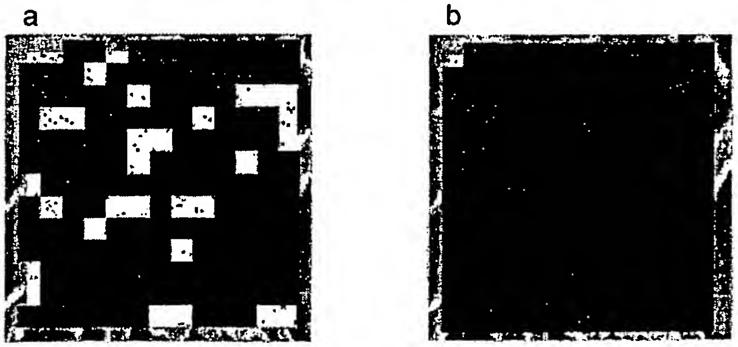

FIG. 12

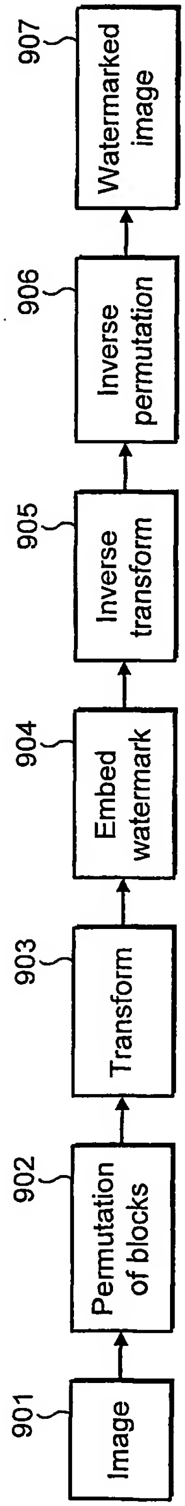
Image
901
Permutation of blocks
902
Transform
903
Embed watermark
904
Inverse transform
905
Inverse permutation
906
Watermarked image
907

FIG. 13

INTERNATIONAL SEARCH REPORT

Int ional Application No
P GB 02/01849

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06T1/00 H04N1/32 H04N7/26 H04N7/24

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06T H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 920 185 A (MATSUSHITA ELECTRIC IND CO LTD) 2 June 1999 (1999-06-02)	1,3-10, 12-15
Y	paragraph ‘0171! - paragraph ‘0222! figures 1-10	2,11
Y	KUNDUR D ET AL: "Digital watermarking using multiresolution wavelet decomposition" ACOUSTICS, SPEECH AND SIGNAL PROCESSING, 1998. PROCEEDINGS OF THE 1998 IEEE INTERNATIONAL CONFERENCE ON SEATTLE, WA, USA 12-15 MAY 1998, NEW YORK, NY, USA,IEEE, US, 12 May 1998 (1998-05-12), pages 2969-2972, XP010279459 ISBN: 0-7803-4428-6 cited in the application the whole document	2

-/--

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

23 July 2002

Date of mailing of the international search report

31/07/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL – 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Moorhouse, D

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 02/01849

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
Y	KUNDUR D ET AL: "DIGITAL WATERMARKING FOR TELLTALE TAMPER PROOFING AND AUTHENTICATION" PROCEEDINGS OF THE IEEE, IEEE. NEW YORK, US, vol. 87, no. 7, July 1999 (1999-07), pages 1167-1180, XP000914459 ISSN: 0018-9219 the whole document	2,11

INTERNATIONAL SEARCH REPORT

Int ıonal Application No
PCı/GB 02/01849

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0920185 A	02-06-1999	JP 11196262 A EP 0920185 A2	21-07-1999 02-06-1999